

TOPS @ Berkeley Lab

Esmond G. Ng

(EGNg@lbl.gov)

Lawrence Berkeley National Laboratory

- ❑ TOPS = Terascale Optimal PDE Simulations.
- ❑ Funded by the Office of Advanced Scientific Computing Research (OASCR) of the Office of Science in the U.S. Department of Energy.
- ❑ Launched in July 2001 under the Scientific Discovery Through Advanced Computing (SciDAC) Program.
 - One of three applied math Integrated Software Infrastructure Centers (ISI C's).
 - One of 51 SciDAC projects in the Office of Science.

- ❑ Solver technology.
- ❑ Not just algorithms, but also vertically integrated software suites.
- ❑ Portable, scalable, extensible, tunable implementations.
- ❑ Starring ARPACK, Hypre, PETSc, SuperLU, and ScaLAPACK, among other existing packages.
- ❑ Motivated by representative applications, intended for many others.

- ❑ Three driving SciDAC science applications in the original plan:
 - LBNL/SLAC-led “Advanced Computing for 21st Century Accelerator Science and Technology”.
 - ORNL-led “Shedding New Light on Exploding Stars: Terascale Simulations of Neutrino-Driven SuperNovae and Their NucleoSynthesis”.
 - PPPL-led “Extended Magnetohydrodynamic Modeling”.
- ❑ Many more application partners now.
 - QCD, chemistry, ...


- ❑ Three DOE laboratories ...
 - Argonne National Laboratory (Jorge Moré, Barry Smith)
 - Lawrence Berkeley National Laboratory (Esmond Ng)
 - Lawrence Livermore National Laboratory (Robert Falgout)
- ❑ Seven universities ...
 - Carnegie Mellon University
 - Columbia University (David Keyes – lead PI)
 - New York University
 - Old Dominion University
 - University of California, Berkeley
 - University of Colorado, Boulder
 - University of Tennessee, Knoxville

□ Include:

- Linear system solvers: $Ax = b$
- Nonlinear implicit solvers: $F(x) = 0$
- Adaptive time integrators for stiff systems: $f(x', x, t) = 0$
- Optimizers: $\min_u \phi(x, u)$ s.t. $F(x, u) = 0$
- Eigenvalue solvers: $Ax = \lambda Bx$

□ Software integration.

□ Performance optimization.



→ Indicates dependence

□ Areas:

- Linear Equations Solvers: $Ax = b$
 - Sparse direct methods
 - Preconditioning techniques for iterative methods
- Eigenvalue Solvers: $Ax = \lambda Bx$

□ Members:

Staff

Parry Husbands

Sherry Li

Osni Marques

Esmond Ng

Chao Yang

Postdocs

Laura Petrescu

Ali Pinar

Visiting Scientist

Weiguo Gao

- ❑ SLAC's Electromagnetic Systems Simulations in the Accelerator Science and Technology SciDAC Project.
 - Linear Algebra – large-scale sparse eigensolvers, sparse linear equations solvers (LBNL, Stanford, SLAC).
 - Load Balancing – improving performance and scalability (LBNL, Sandia, SLAC).


- ❑ Center for Extended Magnetohydrodynamic Modeling.
 - Linear Algebra – solution of large sparse ill-conditioned linear systems (LBNL, Univ. of Wisconsin).

- Modeling of accelerator structures requires the solution of the Maxwell equations.

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} ; \nabla \times \mathbf{H} = \frac{\partial \mathbf{D}}{\partial t}$$


$$\nabla \cdot \mathbf{D} = 0 ; \nabla \cdot \mathbf{B} = 0$$

$$\mathbf{D} = \epsilon \mathbf{E} ; \mathbf{B} = \mu \mathbf{H}$$



- Finite element discretization in frequency domain leads to a large sparse generalized eigenvalue problem.


$$\mathbf{K} \mathbf{x} = \lambda \mathbf{M} \mathbf{x} , \quad \mathbf{K} \geq 0 ; \mathbf{M} > 0$$



- ❑ Design of accelerator structures.
 - Modeling of a single accelerator cell suffices.
 - Relatively small eigenvalue problem.
 - There is an optimization problem here ...
 - But need fast and reliable eigensolvers at every iteration.




- ❑ Understanding the wake field requires the modeling of the full structure.
 - Need to compute a large number of frequency modes.



- ❑ 3-D structures \Rightarrow large matrices.
 - Need very accurate interior eigenvalues that have relatively small magnitudes.
 - Eigenvalues are tightly clustered.
 - When losses in structures are considered, the problems will become complex symmetric.
- ❑ Omega3P has been able to compute eigen modes of large accelerator structures with large number of DOF's (without losses).

Spectral Distribution



- Parallel shift-invert Lanczos algorithm.

- Ideal for computing interior and clustered eigenvalues.


$$K x = I M x \rightarrow M(K - sM)^{-1} M x = m M x$$

- Need solution of sparse linear systems.

- SLAC: **inexact solution + Newton-type correction.**


- LBNL: **exact shift-invert Lanczos.**

- Require complete factorizations of (sparse) matrices.
- Exploit work on sparse direct solvers in TOPS.
- Combine **SuperLU_DIST** with **PARPACK** to obtain a parallel implementation of a shift-invert Lanczos eigensolver.
- Enable accurate calculation of eigenvalues, allow verification of other eigensolvers, and provide a baseline for comparisons.




- ❑ SuperLU and SuperLU_Dist.
 - Direct solution of sparse linear system $Ax = b$.
 - Efficient, high-performance, portable implementations on modern computer architectures.
 - Support real and complex matrices, fill-reducing orderings, equilibration, numerical pivoting, condition estimation, iterative refinement, and error bounds.

Quick Tour of Results



dds47 Performance (16 eigenvalues)

On NERSC IBM SP



□ Accomplishments:

- TOPS/Exact Shift-invert Lanczos and Omega3P/Inexact Shift-invert Lanczos produce the same eigenvalues.
- ESIL is faster than ISIL, but requires more memory.
 - NERSC IBM SP has >6 terabytes of real memory.
- Exact shift-invert Lanczos has been integrated into Omega3P as a run-time option.


□ The accelerator application helps motivate new developments & improvements in SuperLU.

- Accommodate distributed input matrices.
- Parallel symbolic factorization (in progress).
- Improve triangular solution (in progress).


- ❑ Largest eigensystem solved by LBNL team:
 - 7.5M DOFs, 304M nonzeros
 - 6,260M nonzeros in factors, requiring 2.5 hours to compute 10 eigenvalues close to a single shift using 24 processors (on 3 nodes) on NERSC IBM SP.

- ❑ Largest eigensystem solved by SLAC team:
 - > 90M DOFs with a sparse matrix factorization.

- ❑ NI MROD, developed at the Center for Extended Magnetohydrodynamic Modeling (CEMM), is a parallel simulation code for studying the nonlinear macroscopic electromagnetic dynamics in fusion plasmas.
 - Have difficulties in convergence when iterative methods are used in solving linear systems.
- ❑ Joint work between CEMM and TOPS has led to a performance improvement in NI MROD by a factor of 5-10 on NERSC IBM SP, "equivalent of 3-5 years progress in computing hardware" (Dalton Schnack of SAI C).




- ❑ Parallel SuperLU has been incorporated into NI MROD as an alternative linear solver.
 - In updating physical fields, SuperLU is used as a direct solver and is >100x and 64x faster on 1 and 9 processors, respectively.
 - A larger linear system in the time-advance has to be solved by the conjugate gradient method. SuperLU is used to factor a preconditioning matrix. This resulting in a 5-fold improvement in speed.




Load Balancing in EM Simulations


- ❑ Load balancing problem in Tau3P, a time-domain code for electromagnetic simulations.
 - Use of unstructured meshes and refinements lead to matrices for which nonzero entries are not evenly distributed.
 - Make work assignment and load balancing difficult in a parallel setting.
 - SLAC's Tau3P currently uses ParMETIS to partition the domain to minimize communication.




Matrix Sparsity



Matrix Distribution over 14 cpu's




Parallel Speedup




- ❑ Collaboration between SLAC and TOPS (+ Sandia) has resulted in improved performance in Tau3P.
 - Sandia's Zoltan library is implemented to access better partitioning schemes for improved parallel performance over existing ParMETIS tool through reduced communication costs.

8 processor partitioning of a 5-cell RDDS with couplers on NERSC IBM SP


	Tau3P Runtime	Max. Adj. Procs.	Max. Bound. Objects
ParMETIS	288.5 sec	3	585
RCB-1D	218.5 sec	2	3128
RCB-3D	345.6 sec	5	1965



ParMETIS



RCB-1D



RCB-3D

ParMETIS

RCB-1D

Performance results on NERSC IBM SP for a 55-cell structure				
# of processors	ParMETIS run-time	ParMETIS max. adj. procs	RCB-1D run time	RCB-1D max. adj. procs
32	1455.0	4	1236.6	2
64	736.6	4	627.2	2
128	643.0	10	265.1	2
256	360.0	11	129.2	2
512	292.1	14	92.3	4

- ❑ Significant improvement obtained from using RCB-1D over ParMETIS on a 55-cell structure due to the linear nature of the geometry.

- ❑ TOPS/LBNL has contributed to great successes in some SciDAC applications.
 - ❑ Future Plans:
 - Sparse direct solvers.
 - More improvements (e.g., symbolic factorization & triangular solutions) to make sparse direct methods more scalable.
 - Fill-reducing orderings and scheduling issues.
 - Eigenvalue calculations:
 - Other eigen solvers to handle extremely large problems (e.g., multilevel algorithms, Jacobi-Davidson).
 - Incomplete factorization algorithms for iterative methods.
 - Other applications.
- ❑ More information available at <http://www.tops-scidac.org>.